

BOG 4:

OS/Resource Management - Global OS, Composition, Workflow

ASCR Workshop on Extreme Heterogeneity in HPC
23-25 Jan 2018

BOG 4 Contributors

Moderator(s): Ron Brightwell

Scribes: Kevin Pedretti, Andrew Younge

BOGists:

Rob Ross
Lavanya Ramakrishnan
Chris Rossbach
Carlos Maltzahn
Jay Lofstead
Vitus Leung
John Mellor-Crummey
Noah Watkins
Zhiling Lan
Larry Kaplan
Andrew J. Younge
Kevin Pedretti
Jay Lofstead
Craig Vineyard
Paul Peltz

Kenneth Moreland
Pete Beckman
Kamil Iskra
Brian Van Essen
Brian Kocoloski
Douglas Jacobsen
Rebecca Hartman-Baker
Tina Declerck
Anthony Yau
Esam El-Araby
Barry Rountree
Mike Lang
Latchesar Ionkov
Mattan Erez
Suren Byna

Scott Baden
Sriram Krishnamoorthy
Tom Peterka
Valario Pascucci
Andres Marquez
Scott Klasky
Jonathan Beard
Laura Biven
Lucy Nowell
Michael Garland
Pavan Balaji
Peter Kogge
Bruce Jacob
Christine Sweeney

David Richards
Florin Rusu
Franck Cappello
John Wu
Barbara Chapman
Sameer Shende
Abid Malik
Alex Aiken
Terry Jones
Matthew Wolf

Key Research Challenges

- New resource management strategies
 - Support for new types of workloads (interactive, highly dynamic, composed workflows)
 - Managing disaggregated resources (pools of memory, compute, storage)
 - More sophisticated policies for managing resources and applications
 - Hierarchy of resource management (system, rack, node)
- Dynamic resource discovery
 - Mechanisms and interfaces for dynamic resource discovery
 - Communication of resource capabilities and state of hardware
 - Exchange of information between application and Global OS
- Interfaces to the Global OS
 - How the Global OS presents itself as a service
 - Interactions with the RAS system
- Predictability
 - Increasing the predictability of OS services

Complete List of High-Level Topics

- OS/R support for converged memory and storage
- Role of global OS in terms of resource optimization
- Coordination between resource manager and node OS/R
- Mechanisms needed for resource isolation and sharing/coupling
- New types of workloads and GOS functionality needed to support them
- How to manage disaggregated resources (memory, storage, etc.)
- New data abstractions and mechanisms across storage/memory hierarchy
- GOS understanding heterogeneous capabilities, state, resources coming and going
- Co-scheduling workloads for efficient utilization of heterogeneous resources
- Tradeoff between optimizing for resource utilization and performance reproducibility
- Multi-tier resource optimization objectives, granularity is changing. (not just two level)
- Consistency in RAS-type reporting, made more complicated by EH
- Resource discovery and naming, mechanisms for communicating to application. Also need to specify semantics, not simple matching problem
- Autonomic systems and self-healing
- Increased potential of security/privacy issues with more complex hardware
- User interfaces for declarative specification of resource allocation and requirements
- Global network resource management
- Minority report OS, continuously monitoring application and noticing patterns
- Increasing predictability of OS services. Need more expressive QoS specifications.
- What are the kinds of information that an app could be expected to give (in terms of QoS / predictability reqs), and what are benefits that come from those kinds of info?
- Provenance capture of resource management decisions & report to users
- How does GOS present itself to all the other components in system it needs to interact with? What does the GOS API look like?
- How GOS and LOS allocates and schedules in-network operations. Can we adopt technologies from cloud?

Backup

BOG X Capability Targets for Extreme Heterogeneity

BOG X brainstorming and discussion of capabilities that will be needed in the 2025-2035 timeframe to make increasingly heterogeneous hardware technologies useful and productive for science applications.

BOG X Targets for 2030

Target 1:

Target 2:

BOG X Current Status

BOG X survey describing the current status of science, technology, or practice related to this theme (starting with material directly from the FSD).

BOG X Current Capability

Capability 1:

Capability 2:

BOG X Challenge Assessment

Discussion to identify research required to get from where the capabilities are now to where to where they need to be by 2030.

BOG X: list of key research challenges

Challenge X.1

Challenge X.2

BOG X Possible Research Directions Summary

PRD X.1 - direction 1

PRD X.2 - direction 2

PRD X.n - direction n

PRD X.n : Short title of possible research direction

- One paragraph description (3 sentence/bullet)
- Research challenges
 - Metrics for progress
- Potential research approaches and research directions
- How and when will success impact technology?